

Taming Past LTL and Flat Counter Systems^{*}

Stéphane Demri¹, Amit Kumar Dhar², and Arnaud Sangnier²

¹ LSV, CNRS, ENS Cachan, INRIA, France

² LIAFA, Univ Paris Diderot, Sorbonne Paris Cité, CNRS, France

Abstract. Reachability and LTL model-checking problems for flat counter systems are known to be decidable but whereas the reachability problem can be shown in NP, the best known complexity upper bound for the latter problem is made of a tower of several exponentials. Herein, we show that the problem is only NP-complete even if LTL admits past-time operators and arithmetical constraints on counters. Actually, the NP upper bound is shown by adequately combining a new stuttering theorem for Past LTL and the property of small integer solutions for quantifier-free Presburger formulae. Other complexity results are proved, for instance for restricted classes of flat counter systems.

1 Introduction

Flat counter systems. Counter systems are finite-state automata equipped with program variables (counters) interpreted over non-negative integers. They are used in many places like, broadcast protocols [9] and programs with pointers [12] to quote a few examples. But, alongwith their large scope of usability, many problems on general counter systems are known to be undecidable. Indeed, this computational model can simulate Turing machines. Decidability of reachability problems or model-checking problems based on temporal logics, can be regained by considering subclasses of counter systems, see e.g. [14]. An important and natural class of counter systems, in which various practical cases of infinite-state systems (e.g. broadcast protocols [11]) can be modelled, are those with a *flat* control graph, i.e, those where no control state occurs in more than one simple cycle, see e.g. [1,5,11,20]. Decidability results on verifying safety and reachability properties on flat counter systems have been obtained in [5,11,3]. However, so far, such properties have been rarely considered in the framework of any formal specification language (see an exception in [4]). In [7], a class of Presburger counter systems is identified for which the local model checking problem for Presburger-CTL^{*} is shown decidable. These are Presburger counter systems defined over flat control graphs with arcs labelled by adequate Presburger formulae. Even though flatness is clearly a substantial restriction, it is shown in [20] that many classes of counter systems with computable Presburger-definable reachability sets are *flattable*, i.e. there exists a flat unfolding of the counter system with identical reachability sets. Hence, the possibility of flattening a counter system is strongly

^{*} Supported by ANR project REACHARD ANR-11-BS02-001.

related to semilinearity of its reachability set. Moreover, in [4] model-checking relational counter systems over LTL formulae is shown decidable when restricted to flat formulae (their translation into automata leads to flat structures).

Towards the complexity of temporal model-checking flat counter systems. In [7], it is shown that CTL* model-checking over the class of so-called *admissible* counter systems is decidable by reduction into the satisfiability problem for Presburger arithmetic, the decidable first-order theory of natural numbers with addition. Obviously CTL* properties are more expressive than reachability properties but this has a cost. However, for the class of counter systems considered in this paper, this provides a very rough complexity upper bound in 4EXPTIME. Herein, our goal is to revisit standard decidability results for subclasses of counter systems obtained by translation into Presburger arithmetic in order to obtain optimal complexity upper bounds.

Our contributions. In the paper, we establish several computational complexity characterizations of model-checking problems restricted to flat counter systems in the presence of a rich LTL-like specification language with arithmetical constraints and past-time operators. Not only we provide an optimal complexity but also, we believe that our proof technique could be reused for further extensions. Indeed, we combine three proof techniques: the general stuttering theorem [17], the property of small integer solutions of equation systems [2] (this latter technique is used since [24]) and the elimination of disjunctions in guards (see Section 5.2). Let us be a bit more precise.

We extend the stuttering principle established in [17] for LTL (without past-time operators) to Past LTL. The stuttering theorem from [17] for LTL without past-time operators has been used to show that LTL model-checking over *weak* Kripke structures is in NP [16] (weakness corresponds to flatness). It is worth noting that another way to show a similar result would be to eliminate past-time operators thanks to Gabbay’s Separation Theorem [13] (preserving initial equivalence) but the temporal depth of formulae might increase at least exponentially, which is a crucial parameter in our complexity analysis. We show that the model-checking problem restricted to flat counter systems in the presence of LTL with past-time operators is in NP (Theorem 17) by combining the above-mentioned proof techniques. Apart from the use of the general stuttering theorem (Theorem 3), we take advantage of the other properties stated for instance in Lemma 12 (characterization of runs by quantifier-free Presburger formulae) and Theorem 14 (elimination of disjunctions in guards preserving flatness). In the paper, complexity results for fragments/subproblems are also considered. For instance, we get a sharp lower bound since we establish that the model-checking problem on path schemas (a fundamental structure in flat counter systems) with only 2 loops is already NP-hard (see Lemma 11). A summary table can be found in Section 6.

Omitted proofs can be found in [6].

2 Flat Counter Systems and its LTL Dialect

We write \mathbb{N} [resp. \mathbb{Z}] to denote the set of natural numbers [resp. integers] and $[i, j]$ to denote $\{k \in \mathbb{Z} : i \leq k \text{ and } k \leq j\}$. For $\mathbf{v} \in \mathbb{Z}^n$, $\mathbf{v}[i]$ denotes the i^{th} element of \mathbf{v} for every $i \in [1, n]$. For some n -ary tuple t , we write $\pi_j(t)$ to denote the j^{th} element of t ($j \leq n$). In the sequel, integers are encoded with a binary representation. For a finite alphabet Σ , Σ^* represents the set of finite words over Σ , Σ^+ the set of finite non-empty words over Σ and Σ^ω the set of ω -words over Σ . For a finite word $w = a_1 \dots a_k$ over Σ , we write $\text{len}(w)$ to denote its *length* k . For $0 \leq i < \text{len}(w)$, $w(i)$ represents the $(i + 1)$ -th letter of the word, here a_{i+1} .

2.1 Counter Systems

Let $\mathbf{C} = \{x_1, x_2, \dots\}$ be a countably infinite set of *counters* (variables interpreted over non-negative integers) and $\text{AT} = \{p_1, p_2, \dots\}$ be a countable infinite set of propositional variables (abstract properties about program points). We write \mathbf{C}_n to denote $\{x_1, x_2, \dots, x_n\}$. The set $\mathbf{G}(\mathbf{C}_n)$ of *guards* (arithmetical constraints on counters in \mathbf{C}_n) is defined inductively as follows: $\mathbf{t} ::= a.x \mid \mathbf{t} + \mathbf{t}$ and $\mathbf{g} ::= \mathbf{t} \sim b \mid \mathbf{g} \wedge \mathbf{g} \mid \mathbf{g} \vee \mathbf{g}$, where $x \in \mathbf{C}_n$, $a \in \mathbb{Z}$, $b \in \mathbb{Z}$ and $\sim \in \{=, \leq, \geq, <, >\}$. Such guards are closed under negations (but negation is not part of the logical connectives) and the truth constants \top and \perp can be easily defined too. Given $\mathbf{g} \in \mathbf{G}(\mathbf{C}_n)$ and a vector $\mathbf{v} \in \mathbb{N}^n$, we say that \mathbf{v} satisfies \mathbf{g} , written $\mathbf{v} \models \mathbf{g}$, if the formula obtained by replacing each x_i by $\mathbf{v}[i]$ holds.


Definition 1 (Counter system). For $n \geq 1$, a counter system S is a tuple $\langle Q, \mathbf{C}_n, \Delta, \mathbf{l} \rangle$ where Q is a finite set of control states, $\mathbf{l} : Q \rightarrow 2^{\text{AT}}$ is a labelling function and $\Delta \subseteq Q \times \mathbf{G}(\mathbf{C}_n) \times \mathbb{Z}^n \times Q$ is a finite set of edges labeled by guards and updates of the counter values (transitions).

For $\delta = (q, \mathbf{g}, \mathbf{u}, q')$ in Δ , we use the following notations $\text{source}(\delta) = q$, $\text{target}(\delta) = q'$, $\text{guard}(\delta) = \mathbf{g}$ and $\text{update}(\delta) = \mathbf{u}$. As usual, to a counter system $S = \langle Q, \mathbf{C}_n, \Delta, \mathbf{l} \rangle$, we associate a labeled transition system $TS(S) = \langle C, \rightarrow \rangle$ where $C = Q \times \mathbb{N}^n$ is the set of *configurations* and $\rightarrow \subseteq C \times \Delta \times C$ is the *transition relation* defined by: $\langle (q, \mathbf{v}), \delta, (q', \mathbf{v}') \rangle \in \rightarrow$ (also written $(q, \mathbf{v}) \xrightarrow{\delta} (q', \mathbf{v}')$) iff $q = \text{source}(\delta)$, $q' = \text{target}(\delta)$, $\mathbf{v} \models \text{guard}(\delta)$ and $\mathbf{v}' = \mathbf{v} + \text{update}(\delta)$. In such a transition system, the counter values are non-negative since $C = Q \times \mathbb{N}^n$. We extend the transition relation \rightarrow to finite words of transitions in Δ^+ as follows. For each $w = \delta_1 \delta_2 \dots \delta_\alpha \in \Delta^+$, we have $\langle (q, \mathbf{v}), w \rangle \rightarrow \langle (q', \mathbf{v}') \rangle$ if there are $c_0, c_1, \dots, c_{\alpha+1} \in C$ such that $c_i \xrightarrow{\delta_i} c_{i+1}$ for all $i \in [0, \alpha]$, $c_0 = (q, \mathbf{v})$ and $c_{\alpha+1} = \langle (q', \mathbf{v}') \rangle$. We say that an ω -word $w \in \Delta^\omega$ is *fireable* in S from a configuration $c_0 \in Q \times \mathbb{N}^n$ if for all finite prefixes w' of w there exists a configuration $c \in Q \times \mathbb{N}^n$ such that $c_0 \xrightarrow{w'} c$. We write $\text{lab}(c_0)$ to denote the set of ω -words (*labels*) which are fireable from c_0 in S .

Given a configuration $c_0 \in Q \times \mathbb{N}^n$, a *run* ρ starting from c_0 in S is an infinite path in the associated transition system $TS(S)$ denoted as: $\rho := c_0 \xrightarrow{\delta_0} \dots$

$\dots \xrightarrow{\delta_{\alpha-1}} c_\alpha \xrightarrow{\delta_\alpha} \dots$ where $c_i \in Q \times \mathbb{N}^n$ and $\delta_i \in \Delta$ for all $i \in \mathbb{N}$. Let $lab(\rho)$ be the ω -word $\delta_0\delta_1\dots$ associated to the run ρ . Note that by definition we have $lab(\rho) \in lab(c_0)$. When E is an ω -regular expression over the finite alphabet Δ and c_0 is an initial configuration, $lab(E, c_0)$ is defined as the set of labels of infinite runs ρ starting at c_0 such that $lab(\rho)$ belongs to the language defined by E . So $lab(E, c_0) \subseteq lab(c_0)$.

We say that a counter system is *flat* if every node in the underlying graph belongs to at most one simple cycle (a cycle being simple if no edge is repeated twice in it) [5]. In a flat counter system, simple cycles can be organized as a DAG where two simple cycles are in the relation whenever there is path between a node of the first cycle and a node of the second cycle. We denote by \mathcal{CFS} the class of flat counter systems.



On the left, we present the control graph of a flat counter system (guards and updates are omitted). A *Kripke structure* S is a tuple $\langle Q, \Delta, \mathbf{l} \rangle$ where $\Delta \subseteq Q \times Q$ and \mathbf{l} is labelling. It can be viewed as a degenerate form of counter systems without counters (in the sequel, we take the freedom to see them as counter systems). All standard notions on counter systems naturally apply to Kripke structures too (configuration, run, flatness, etc.). In the sequel, we shall also investigate the complexity of model-checking problems on flat Kripke structures (such a class is denoted by \mathcal{KFS}).

2.2 Linear Temporal Logic with Past and Arithmetical Constraints

Model-checking problem for Past LTL over finite state systems is known to be PSPACE-complete. In spite of this nice feature, a propositional variable p only represents an abstract property about the current configuration of the system. A more satisfactory solution is to include in the logical language the possibility to express directly constraints between variables of the program, whence giving up the standard abstraction made with propositional variables. We define below a version of LTL dedicated to counter systems in which the atomic formulae are linear constraints; this is analogous to the use of concrete domains in description logics [21]. Note that capacity constraints from [8] are arithmetical constraints different from those defined below. Formulae of PLTL[C] are defined from $\phi ::= p \mid \mathbf{g} \mid \neg\phi \mid \phi \wedge \phi \mid \phi \vee \phi \mid \mathbf{X}\phi \mid \phi\mathbf{U}\phi \mid \mathbf{X}^{-1}\phi \mid \phi\mathbf{S}\phi$ where $p \in \text{AT}$ and $\mathbf{g} \in \mathbf{G}(\mathbb{C}_n)$ for some n . We may use the standard abbreviations $\mathbf{F}, \mathbf{G}, \mathbf{G}^{-1}$ etc. For instance, the formula $\mathbf{GF}(x_1 + 2 \geq x_2)$ states that infinitely often the value of counter 1 plus 2 is greater than the value of counter 2. The past-time operators \mathbf{S} and \mathbf{X}^{-1} do not add expressive power to the logic itself, but it is known that it helps a lot to express properties succinctly, see e.g. [19,18]. The temporal depth of ϕ , written $td(\phi)$, is defined as the maximal number of imbrications of temporal operators in ϕ . Restriction of PLTL[C] to atomic formulae from AT only is written PLTL[\emptyset], standard version of LTL with past-time operators. Models of PLTL[C] are essentially abstractions of runs from counter systems,

i.e. ω -sequences $\sigma : \mathbb{N} \rightarrow 2^{\text{AT}} \times \mathbb{N}^{\text{C}}$. Given a model σ and a position $i \in \mathbb{N}$, the satisfaction relation \models for PLTL[C] is defined as follows (other cases can be defined similarly, see e.g. [18]):

- $\sigma, i \models p \stackrel{\text{def}}{\iff} p \in \pi_1(\sigma(i)), \sigma, i \models \mathbf{g} \stackrel{\text{def}}{\iff} \mathbf{v}_i \models \mathbf{g}$ where $\mathbf{v}_i[j] \stackrel{\text{def}}{=} \pi_2(\sigma(i))(x_j)$,
- $\sigma, i \models \mathbf{X}\phi \stackrel{\text{def}}{\iff} \sigma, i+1 \models \phi$,
- $\sigma, i \models \phi_1 \mathbf{S} \phi_2 \stackrel{\text{def}}{\iff} \sigma, j \models \phi_2$ for some $0 \leq j \leq i$ s.t. $\sigma, k \models \phi_1, \forall j < k \leq i$.

Given $\langle Q, \mathbf{C}_n, \Delta, \mathbf{I} \rangle$ and a run $\rho := \langle q_0, \mathbf{v}_0 \rangle \xrightarrow{\delta_0} \dots \xrightarrow{\delta_{p-1}} \langle q_p, \mathbf{v}_p \rangle \xrightarrow{\delta_p} \dots$, we consider the model $\sigma_\rho : \mathbb{N} \rightarrow 2^{\text{AT}} \times \mathbb{N}^{\text{C}}$ such that $\pi_1(\sigma_\rho(i)) \stackrel{\text{def}}{=} \mathbf{I}(q_i)$ and $\pi_2(\sigma_\rho(i))(x_j) \stackrel{\text{def}}{=} \mathbf{v}_i[j]$ for all $j \in [1, n]$ and all $i \in \mathbb{N}$. Note that $\pi_2(\sigma_\rho(i))(x_j)$ is arbitrary for $j \notin [1, n]$. As expected, we extend the satisfaction relation to runs so that $\rho, i \models \phi \stackrel{\text{def}}{\iff} \sigma_\rho, i \models \phi$ whenever ϕ is built from counters in \mathbf{C}_n .

Given a fragment L of PLTL[C] and a class \mathcal{C} of counter systems, we write $\text{MC}(\text{L}, \mathcal{C})$ to denote the existential model checking problem: given $S \in \mathcal{C}$, a configuration c_0 and $\phi \in \text{L}$, does there exist ρ starting from c_0 such that $\rho, 0 \models \phi$? In that case, we write $S, c_0 \models \phi$. It is known that for the full class of counter systems, the model-checking problem is undecidable, see e.g. [22]. Some restrictions, such as flatness, can lead to decidability as shown in [7] but the decision procedure there involves an exponential reduction to Presburger Arithmetic, whence the high complexity.

Theorem 2. [7,16] $\text{MC}(\text{PLTL}[\mathbf{C}], \mathcal{CFS})$ can be solved in 4EXPTIME . $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KFS})$ restricted to formulae with temporal operators \mathbf{U}, \mathbf{X} is NP-complete.

Our main goal is to characterize the complexity of $\text{MC}(\text{PLTL}[\mathbf{C}], \mathcal{CFS})$.

3 Stuttering Theorem for PLTL[\emptyset]

Stuttering of finite words or single letters has been instrumental to show results about the expressive power of PLTL[\emptyset] fragments, see e.g. [23,17]; for instance, PLTL[\emptyset] restricted to the temporal operator \mathbf{U} characterizes the class of formulae defining classes of models invariant under stuttering. This is refined in [17] for PLTL[\emptyset] restricted to \mathbf{U} and \mathbf{X} , by taking into account not only the \mathbf{U} -depth but also the \mathbf{X} -depth of formulae and by introducing a principle of stuttering that involves both letter stuttering and word stuttering. In this section, we establish another substantial generalization that involves PLTL[\emptyset] with past-time temporal operators. Roughly speaking, we show that if $\sigma_1 \mathbf{s}^M \sigma_2, 0 \models \phi$ where $\sigma_1 \mathbf{s}^M \sigma_2$ is a PLTL[\emptyset] model (σ_1, \mathbf{s} being finite words), $\phi \in \text{PLTL}[\emptyset]$, $td(\phi) \leq N$ and $M \geq 2N + 1$, then $\sigma_1 \mathbf{s}^{2N+1} \sigma_2, 0 \models \phi$ (and other related properties). This extends a result without past-time operators [16]. Moreover, this turns out to be a key property (Theorem 3) to establish the NP upper bound even in the presence of counters. Note that Theorem 3 below is interesting for its own sake, independently of our investigation on flat counter systems. By lack of space, we state below the main definitions and result.

Given $M, M', N \in \mathbb{N}$, we write $M \approx_N M'$ iff $\text{Min}(M, N) = \text{Min}(M', N)$. Given $w = w_1 u^M w_2, w' = w_1 u^{M'} w_2 \in \Sigma^\omega$ and $i, i' \in \mathbb{N}$, we define an equivalence relation $\langle w, i \rangle \approx_N \langle w', i' \rangle$ (implicitly parameterized by w_1, w_2 and u) such that $\langle w, i \rangle \approx_N \langle w', i' \rangle$ means that the number of copies of u before position i and the number of copies of u before position i' are related by \approx_N and the same applies for the number of copies after the positions. Moreover, if i and i' occur in the part where u is repeated, then they correspond to identical positions in u . More formally, $\langle w, i \rangle \approx_N \langle w', i' \rangle \stackrel{\text{def}}{\iff} M \approx_{2N} M'$ and one of the conditions holds true: (1) $i, i' < \text{len}(w_1) + N \cdot \text{len}(u)$ and $i = i'$, (2) $i \geq \text{len}(w_1) + (M - N) \cdot \text{len}(u)$, $i' \geq \text{len}(w_1) + (M' - N) \cdot \text{len}(u)$ and $(i - i') = (M - M') \cdot \text{len}(u)$, (3) $\text{len}(w_1) + N \cdot \text{len}(u) \leq i < \text{len}(w_1) + (M - N) \cdot \text{len}(u)$, $\text{len}(w_1) + N \cdot \text{len}(u) \leq i' < \text{len}(w_1) + (M' - N) \cdot \text{len}(u)$ and $|i - i'| = 0 \pmod{\text{len}(u)}$. We state our stuttering theorem for PLTL $[\emptyset]$ that is tailored for our future needs.

Theorem 3 (Stuttering). *Let $\sigma = \sigma_1 \mathbf{s}^M \sigma_2, \sigma' = \sigma_1 \mathbf{s}^{M'} \sigma_2 \in (2^{\text{AT}})^\omega$ and $i, i' \in \mathbb{N}$ such that $N \geq 2$, $M, M' \geq 2N + 1$ and $\langle \sigma, i \rangle \approx_N \langle \sigma', i' \rangle$. Then, for every PLTL $[\emptyset]$ formula ϕ with $\text{td}(\phi) \leq N$, we have $\sigma, i \models \phi$ iff $\sigma', i' \models \phi$.*

Proof. (sketch) The proof is by structural induction on the formula but first we need to establish properties whose proofs can be found in [6]. Let $w = w_1 u^M w_2, w' = w_1 u^{M'} w_2 \in \Sigma^\omega$, $i, i' \in \mathbb{N}$ and $N \geq 2$ such that $M, M' \geq 2N + 1$ and $\langle w, i \rangle \approx_N \langle w', i' \rangle$. We can show the following properties:

- (Claim 1) $\langle w, i \rangle \approx_{N-1} \langle w', i' \rangle$ and $w(i) = w'(i')$.
- (Claim 2) $\langle w, i + 1 \rangle \approx_{N-1} \langle w', i' + 1 \rangle$ and $i, i' > 0$ implies $\langle w, i - 1 \rangle \approx_{N-1} \langle w', i' - 1 \rangle$.
- (Claim 3) For all $j \geq i$, there is $j' \geq i'$ such that $\langle w, j \rangle \approx_{N-1} \langle w', j' \rangle$ and for all $k' \in [i', j' - 1]$, there is $k \in [i, j - 1]$ such that $\langle w, k \rangle \approx_{N-1} \langle w', k' \rangle$.
- (Claim 4) For all $j \leq i$, there is $j' \leq i'$ such that $\langle w, j \rangle \approx_{N-1} \langle w', j' \rangle$ and for all $k' \in [j' - 1, i']$, there is $k \in [j - 1, i]$ such that $\langle w, k \rangle \approx_{N-1} \langle w', k' \rangle$.

By way of example, let us present the induction step for subformulae of the form $\psi_1 \mathbf{U} \psi_2$. We show that $\sigma, i \models \psi_1 \mathbf{U} \psi_2$ implies $\sigma', i' \models \psi_1 \mathbf{U} \psi_2$. Suppose there is $j \geq i$ such that $\sigma, j \models \psi_2$ and for every $k \in [i, j - 1]$, we have $\sigma, k \models \psi_1$. There is $j' \geq i'$ satisfying (Claim 3). Since $\text{td}(\psi_1), \text{td}(\psi_2) \leq N - 1$, by (IH), we have $\sigma', j' \models \psi_2$. Moreover, for every $k' \in [i', j' - 1]$, there is $k \in [i, j - 1]$ such that $\langle w, k \rangle \approx_{N-1} \langle w', k' \rangle$ and by (IH), we have $\sigma', k' \models \psi_1$ for every $k' \in [i', j' - 1]$. Hence, $\sigma', i' \models \psi_1 \mathbf{U} \psi_2$. \square

An alternative proof consists in using Ehrenfeucht-Fraïssé games [10].

4 Fundamental Structures: Minimal Path Schemas

In this section, we introduce the notion of a fundamental structure for flat counter systems, namely a path schema. Indeed, every flat counter system can be decomposed into a finite set of minimal path schemas and there are only an exponential number of them. So, all our nondeterministic algorithms on flat counter systems have a preliminary step that first guesses a minimal path schema.

4.1 Minimal Path Schemas

Let $S = \langle Q, \mathbf{C}_n, \Delta, \mathbf{1} \rangle$ be a flat counter system. A *path segment* p of S is a finite sequence of transitions from Δ such that $\text{target}(p(i)) = \text{source}(p(i+1))$ for all $0 \leq i < \text{len}(p) - 1$. We write $\text{first}(p)$ [resp. $\text{last}(p)$] to denote the first [resp. last] control state of a path segment, in other words $\text{first}(p) = \text{source}(p(0))$ and $\text{last}(p) = \text{target}(p(\text{len}(p) - 1))$. We also write $\text{effect}(p)$ to denote the sum vector $\sum_{0 \leq i < \text{len}(p)} \text{update}(p(i))$ representing the total effect of the updates along the path segment. A path segment p is said to be *simple* if $\text{len}(p) > 0$ and for all $0 \leq i, j < \text{len}(p)$, $p(i) = p(j)$ implies $i = j$ (no repetition of transitions). A *loop* is a simple path segment p such that $\text{first}(p) = \text{last}(p)$. A *path schema* P is an ω -regular expression built over Δ such that its language represents an overapproximation of the set of labels obtained from infinite runs following the transitions of P . A path schema P is of the form $p_1 l_1^+ p_2 l_2^+ \dots p_k l_k^\omega$ where (1) l_1, \dots, l_k are loops and (2) $p_1 l_1 p_2 l_2 \dots p_k l_k$ is a path segment.

We write $\text{len}(P)$ to denote $\text{len}(p_1 l_1 p_2 l_2 \dots p_k l_k)$ and $\text{nbloops}(P)$ as its number k of loops. Let $\mathcal{L}(P)$ denote the set of infinite words in Δ^ω which belong to the language defined by P . Note that some elements of $\mathcal{L}(P)$ may not correspond to any run because of constraints on counter values. Given $w \in \mathcal{L}(P)$, we write $\text{iter}_P(w)$ to denote the unique tuple in $(\mathbb{N} \setminus \{0\})^{k-1}$ such that $w = p_1 l_1^{\text{iter}_P(w)[1]} p_2 l_2^{\text{iter}_P(w)[2]} \dots p_k l_k^\omega$. So, for every $i \in [1, k-1]$, $\text{iter}_P(w)[i]$ is the number of times the loop l_i is taken. Then, for a configuration c_0 , the set $\text{iter}_P(c_0)$ is the set of vectors $\{\text{iter}_P(w) \in (\mathbb{N} \setminus \{0\})^{k-1} \mid w \in \text{lab}(P, c_0)\}$. Finally, we say that a run ρ starting in a configuration c_0 *respects* a path schema P if $\text{lab}(\rho) \in \text{lab}(P, c_0)$ and for such a run, we write $\text{iter}_P(\rho)$ to denote $\text{iter}_P(\text{lab}(\rho))$. Note that by definition, if ρ respects P , then each loop l_i is visited at least once, and the last one infinitely.

So far, a flat counter system may have an infinite set of path schemas. However, we can impose minimality conditions on path schemas without sacrificing completeness. A path schema $p_1 l_1^+ p_2 l_2^+ \dots p_k l_k^\omega$ is *minimal* whenever $p_1 \dots p_k$ is either the empty word or a simple non-loop segment, and l_1, \dots, l_k are loops with disjoint sets of transitions.

Lemma 4. *Given a flat counter system $S = \langle Q, \mathbf{C}_n, \Delta, \mathbf{1} \rangle$, the total number of minimal path schemas of S is finite and is smaller than $\text{card}(\Delta)^{(2 \times \text{card}(\Delta))}$.*

This is a simple consequence of the fact that in a minimal path schema, each transition occurs at most twice. In Figure 1, we present a flat counter system S with a unique counter and one of its minimal path schemas. Each transition δ_i labelled by $+i$ corresponds to a transition with the guard \top and the update value $+i$. The minimal path schema shown in Figure 1 corresponds to the ω -regular expression $\delta_1(\delta_2\delta_3)^+\delta_4\delta_5(\delta_6\delta_5)^\omega$. Note that in the representation of path schemas, a state may occur several times, as it is the case for q_3 (this cannot occur in the representation of counter systems). Minimal path schemas play a crucial role in the sequel. Indeed, given a path schema P , there is a minimal path schema P' such that every run respecting P respects P' too. This can be easily




Fig. 1. A flat counter system and one of its minimal path schemas

shown since whenever a maximal number of copies of a simple loop is identified as a factor of $p_1 l_1 \cdots p_k l_k$, this factor is replaced by the simple loop unless it is already present in the path schema.

Finally, the conditions imposed on the structure of path schemas implies the following corollary which states that the number of minimal path schemas for a given flat counter system is at most exponential in the size of the system (see similar statements in [20]).

Corollary 5. *Given a flat counter system S and a configuration c_0 , there is a finite set of minimal path schemas X of cardinality at most $\text{card}(\Delta)^{(2 \times \text{card}(\Delta))}$ such that $\text{lab}(c_0) = \text{lab}(\bigcup_{P \in X} P, c_0)$.*

4.2 Complexity Results

We write \mathcal{CPS} [resp. \mathcal{KPS}] to denote the class of path schemas from counter systems [resp. the class of path schemas from Kripke structures]. As a preliminary step, we consider the problem $\text{MC}(\text{PLTL}[\emptyset], \mathcal{CPS})$ that takes as inputs a path schema P in \mathcal{CPS} , and $\phi \in \text{PLTL}[\emptyset]$ and asks whether there is a run respecting P that satisfies ϕ . Let ρ and ρ' be runs respecting P . For $\alpha \geq 0$, we write $\rho \equiv_\alpha \rho' \stackrel{\text{def}}{=} \forall i \in [1, \text{nbloops}(P) - 1]$, we have $\text{Min}(\text{iter}_P(\rho)[i], \alpha) = \text{Min}(\text{iter}_P(\rho')[i], \alpha)$. We state below a result concerning the runs of flat counter systems when respecting the same path schema.

Proposition 6. *Let S be a flat counter system, P be a path schema, and $\phi \in \text{PLTL}[\emptyset]$. For all runs ρ and ρ' respecting P such that $\rho \equiv_{2\text{id}(\phi)+5} \rho'$, we have $\rho, 0 \models \phi$ iff $\rho', 0 \models \phi$.*

This property can be proved by applying Theorem 3 repeatedly in order to get rid of the unwanted iterations of the loops. Our algorithm for $\text{MC}(\text{PLTL}[\emptyset], \mathcal{CPS})$ takes advantage of a result from [18] for model-checking ultimately periodic models with formulae from Past LTL. An *ultimately periodic path* is an infinite word in Δ^ω of the form uv^ω where uv is a path segment and consequently $\text{first}(v) = \text{last}(v)$. According to [18], given an ultimately periodic path w , and a formula $\phi \in \text{PLTL}[\emptyset]$, the problem of checking whether there exists a run ρ such that $\text{lab}(\rho) = w$ and $\rho, 0 \models \phi$ is in PTIME (a tighter bound of NC can be obtained by combining results from [15] and Theorem 3).

Lemma 7. $\text{MC}(\text{PLTL}[\emptyset], \mathcal{CPS})$ is in NP.

The proof is a consequence of Proposition 6 and [18]. Indeed, given $\phi \in \text{PLTL}[\emptyset]$ and $P = p_1 l_1^+ p_2 l_2^+ \dots p_k l_k^\omega$, first guess $\mathbf{m} \in [1, 2td(\phi) + 5]^{k-1}$ and check whether $\rho, 0 \models \phi$ where ρ is the obvious ultimately periodic word such that $\text{lab}(\rho) = p_1 l_1^{\mathbf{m}[1]} p_2 l_2^{\mathbf{m}[2]} \dots p_k l_k^\omega$. Since \mathbf{m} is of polynomial size and $\rho, 0 \models \phi$ can be checked in polynomial time by [18], we get the NP upper bound.

From [16], we have the lower bound for $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KPS})$.

Lemma 8. [16] $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KPS})$ is NP-hard even if restricted to \mathbf{X} and \mathbf{F} .

For a fixed $n \in \mathbb{N}$, we write $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KPS}(n))$ to denote the restriction of $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KPS})$ to path schemas with at most n loops. When n is fixed, the number of ultimately periodic paths w in $\mathcal{L}(P)$ such that each loop (except the last one) is visited at most $2td(\phi) + 5$ times is bounded by $(2td(\phi) + 5)^n$, which is polynomial in the size of the input (because n is fixed).

Theorem 9. $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KPS})$ is NP-complete.

Given a fixed $n \in \mathbb{N}$, $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KPS}(n))$ is in PTIME.

Note that it can be proved that $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KPS}(n))$ is in NC, hence giving a tighter upper bound for the problem. This can be obtained by observing that we can run the NC algorithm for model checking $\text{PLTL}[\emptyset]$ over ultimately periodic paths parallelly on $(2td(\phi) + 5)^n$ (polynomially many) different paths.

Now, we present how to solve $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KFS})$ using Lemma 7. From Lemma 4, we know that the number of minimal path schemas in a flat Kripke structure $S = \langle Q, \Delta, \mathbf{1} \rangle$ is finite and the length of a minimal path schema is at most $2 \times \text{card}(\Delta)$. Hence, for solving the model-checking problem for a state q and a $\text{PLTL}[\emptyset]$ formula ϕ , a possible algorithm consists in choosing nondeterministically a minimal path schema P starting at q and then apply the algorithm used to establish Lemma 7. This new algorithm would be in NP. Furthermore, thanks to Corollary 5, we know that if there exists a run ρ of S such that $\rho, 0 \models \phi$ then there exists a minimal path schema P such that ρ respects P . Consequently there is an algorithm in NP to solve $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KFS})$.

Theorem 10. $\text{MC}(\text{PLTL}[\emptyset], \mathcal{KFS})$ is NP-complete.

NP-hardness can be established as a variant of the proof of Lemma 8.

Similarly, $\mathcal{CPS}(k)$ denotes the class of path schemas obtained from flat counter systems with number of loops bounded by k .

Lemma 11. For $k \geq 2$, $\text{MC}(\text{PLTL}[\mathbf{C}], \mathcal{CPS}(k))$ is NP-hard.

The proof by reduction from SAT and it is less straightforward than the proof for Lemma 8 or the reduction presented in [16] when path schemas are involved. Indeed, we cannot encode the nondeterminism in the structure itself and the structure has only a constant number of loops. Actually, we cannot use a separate loop for each counter; the reduction is done by encoding the nondeterminism in the (possibly exponential) number of times a single loop is taken, and then using its binary encoding as an assignment for the propositional variables. Hence, the reduction uses in an essential way the counter values and the arithmetical constraints in the formula. By contrast, $\text{MC}(\text{PLTL}[\mathbf{C}], \mathcal{CPS}(1))$ can be shown in PTIME.

5 Model-checking PLTL[C] over Flat Counter Systems

In this section, we provide a nondeterministic polynomial-time algorithm to solve MC(PLTL[C], \mathcal{CFS}) (see Algorithm 1). To do so, we combine Theorem 3 with small solutions of constraint systems.

5.1 Characterizing Runs by System of Equations

In this section, we show how to build a system of equations from a path schema P and a configuration c_0 such that the system of equations encodes the set of all runs respecting P from c_0 . This can be done for path schemas without disjunctions in guards that satisfy an additional *validity* property. A path schema $P = p_1 l_1^+ p_2 l_2^+ \dots p_k l_k^\omega$ is *valid* whenever $effect(l_k)[i] \geq 0$ for every $i \in [1, n]$ (see Section 4 for the definition of $effect(l_k)$) and if all the guards in transitions in l_k are conjunctions of atomic guards, then for each guard occurring in the loop l_k of the form $\sum_i a_i x_i \sim b$ with $\sim \in \{\leq, <\}$ [resp. with $\sim \in \{=\}$, with $\sim \in \{\geq, >\}$], we have $\sum_i a_i \times effect(l_k)[i] \leq 0$ [resp. $\sum_i a_i \times effect(l_k)[i] = 0$, $\sum_i a_i \times effect(l_k)[i] \geq 0$]. It is easy to check that these conditions are necessary to visit the last loop l_k infinitely. More specifically, if a path schema is not valid, then no infinite run can respect it. Moreover, given a path schema, one can decide in polynomial time whether it is valid.

Now, let us consider a (not necessarily minimal) valid path schema $P = p_1 l_1^+ p_2 l_2^+ \dots p_k l_k^\omega$ ($k \geq 1$) obtained from a flat counter system S such that all the guards on transitions are conjunctions of atomic guards of the form $\sum_i a_i x_i \sim b$ where $a_i \in \mathbb{Z}$, $b \in \mathbb{Z}$ and $\sim \in \{=, \leq, \geq, <, >\}$. Hence, disjunctions are *disallowed* in guards. The goal of this section (see Lemma 12 below) is to characterize the set $iter_P(c_0) \subseteq \mathbb{N}^{k-1}$ for some configuration c_0 as the set of solutions of a constraint system. For each loop l_i , we introduce a variable y_i , whence the number of variables of the system/formula is precisely $k - 1$. A *constraint system* \mathcal{E} over the set of variables $\{y_1, \dots, y_n\}$ is a quantifier-free Presburger formula built over $\{y_1, \dots, y_n\}$ as a conjunction of atomic constraints of the form $\sum_i a_i y_i \sim b$ where $a_i, b \in \mathbb{Z}$ and $\sim \in \{=, \leq, \geq, <, >\}$. Conjunctions of atomic counter constraints and constraint systems are essentially the same objects but the distinction allows to emphasize the different purposes: guard on counters in operational models and symbolic representation of sets of tuples.

Lemma 12. *Let $S = \langle Q, C_n, \Delta, \mathbf{1} \rangle$ be a flat counter system without disjunctions in guards, P be a valid path schema and c_0 be a configuration. One can compute in polynomial time a constraint system \mathcal{E} such that the set of solutions of \mathcal{E} is equal to $iter_P(c_0)$, \mathcal{E} has $\text{nbloops}(P) - 1$ variables, \mathcal{E} has at most $\text{len}(P) \times 2 \times \text{size}(S)^2$ conjuncts and the greatest absolute value from constants in \mathcal{E} is bounded by $n \times \text{nbloops}(P) \times K^4 \times \text{len}(P)^3$ where K is the greatest absolute value for constants occurring in S .*

5.2 Elimination of Arithmetical Constraints and Disjunctions

As stated in Lemma 12, the procedure for characterizing infinite runs in a counter system by a system of equations works only for a flat counter system with no

disjunction in guards (convexity of guards is essential). In this section, we show how to obtain such a system from a general flat counter system. Given a flat counter system $S = \langle Q, \mathcal{C}_n, \Delta, \mathbf{1} \rangle$, a configuration $c_0 = \langle q_0, \mathbf{v}_0 \rangle$ and a minimal path schema P starting from the configuration c_0 , we show that it is possible to build a finite set Y_P of path schemas such that (1) each path schema in Y_P has transitions without disjunctions in guards, (2) existence of a run ρ respecting P is equivalent to the existence of a path schema in Y_P having a run similar to ρ respecting it and (3) each path schema in Y_P is obtained from P by unfolding loops so that the terms in each loop satisfy the same atomic guards. Note that disjunctions could be easily eliminated at the cost of adding new transitions between states but this type of transformation may easily destroy flatness. Hence, the necessity to present a more sophisticated elimination procedure

We first introduce a few definitions. A (syntactic) *resource* R is a triple $\langle X, T, B \rangle$ such that X is a finite set of propositional variables, T is a finite set of terms \mathfrak{t} appearing in some guards of the form $\mathfrak{t} \sim b$ (with $b \in \mathbb{Z}$) and B is a finite set of integers. We say that a resource $R = \langle X, T, B \rangle$ is *coherent* with a counter system S [resp. with a path schema P] if B contains all the constants b occurring in guards of S [resp. of P] of the form $\mathfrak{t} \sim b$ and T contains all the terms \mathfrak{t} occurring in guards of S [resp. of P] of the form $\mathfrak{t} \sim b$. The resource R is coherent with a formula $\phi \in \text{PLTL}[\mathcal{C}]$, whenever the atomic formulae of ϕ are either of the form $p \in X$ or $\mathfrak{t} \sim b$ with $\mathfrak{t} \in T$ and $b \in B$. In the sequel, we assume that the considered resource is always coherent with S .

Assuming that $B = \{b_1, \dots, b_m\}$ with $b_1 < \dots < b_m$, we write I to denote the finite set of intervals $I = \{(-\infty, b_1 - 1], [b_1, b_1], [b_1 + 1, b_2 - 1], [b_2, b_2], \dots, [b_m, b_m], [b_m + 1, \infty)\}$. Note that I contains exactly $2m + 1$ intervals. A *term map* \mathbf{m} is a map $\mathbf{m} : T \rightarrow I$ that abstracts term values. A *footprint* is an abstraction of a model for $\text{PLTL}[\mathcal{C}]$ restricted to elements from the resource R : it is of the form $\text{ft} : \mathbb{N} \rightarrow 2^X \times I^T$ where I is the set of intervals built from B . The satisfaction relation \models involving models or runs can be adapted to footprints as follows (formulae and footprints are from the same resource):

- $\text{ft}, i \models_{\text{symb}} p \stackrel{\text{def}}{\iff} p \in \pi_1(\text{ft}(i)); \text{ft}, i \models_{\text{symb}} \mathfrak{t} \geq b \stackrel{\text{def}}{\iff} \pi_2(\text{ft}(i))(\mathfrak{t}) \subseteq [b, +\infty),$
- $\text{ft}, i \models_{\text{symb}} \mathfrak{t} \leq b \stackrel{\text{def}}{\iff} \pi_2(\text{ft}(i))(\mathfrak{t}) \subseteq (-\infty, +b],$
- $\text{ft}, i \models_{\text{symb}} X\phi \stackrel{\text{def}}{\iff} \text{ft}, i + 1 \models_{\text{symb}} \phi,$
- $\text{ft}, i \models_{\text{symb}} \phi U \psi \stackrel{\text{def}}{\iff} \exists j \geq i \text{ s.t. } \text{ft}, j \models_{\text{symb}} \psi \text{ and } \forall j' \in [i, j - 1], \text{ft}, j' \models_{\text{symb}} \phi.$

We omit the other obvious clauses. \models_{symb} is the satisfaction relation for Past LTL when arithmetical constraints are understood as abstract propositions. Let $R = \langle X, T, B \rangle$ be a resource and $\rho = \langle q_0, \mathbf{v}_0 \rangle, \langle q_1, \mathbf{v}_1 \rangle \dots$ be an infinite run of S . The *footprint* of ρ with respect to R is the footprint $\text{ft}(\rho)$ such that for $i \geq 0$, we have $\text{ft}(\rho)(i) \stackrel{\text{def}}{=} \langle \mathbf{1}(q_i) \cap X, \mathbf{m}_i \rangle$ where for every term $\mathfrak{t} = \sum_j a_j x_j \in T$, we have $\sum_j a_j \mathbf{v}_i[j] \in \mathbf{m}_i(\mathfrak{t})$. Note that $\sum_j a_j \mathbf{v}_i[j]$ belongs to a unique element of I since I is a partition of \mathbb{Z} . Hence, this definition makes sense. Lemma 13 below roughly states that satisfaction of a formula on a run can be checked symbolically from the footprint (this is useful for the correctness of forthcoming Algorithm 1).

Lemma 13. *Let ϕ be in PLTL[C], $R = \langle X, T, B \rangle$ be coherent with ϕ , $\rho = \langle q_0, \mathbf{v}_0 \rangle, \langle q_1, \mathbf{v}_1 \rangle \dots$ be an infinite run and $i \geq 0$. (I) Then $\rho, i \models \phi$ iff $\text{ft}(\rho), i \models_{\text{symbol}} \phi$. (II) If ρ' is an infinite run s.t. $\text{ft}(\rho) = \text{ft}(\rho')$, then $\rho, i \models \phi$ iff $\rho', i \models \phi$.*

In [6] we explain in details how to build a set Y_P of path schemas without disjunctions from a minimal path schema P , an initial configuration $\langle q_0, \mathbf{v}_0 \rangle$ and a resource R . The main idea of this construction consists in adding to the control states of path schemas some information on the intervals to which belongs each term of T . In fact, in the transitions appearing in path schemas of Y_P the states belong to $Q' = Q \times I^T$. Before stating the properties of Y_P , we introduce some notations. Given $\mathbf{t} = \sum_j a_j x_j \in T$, $\mathbf{u} \in \mathbb{Z}^n$ and a term map \mathbf{m} , we write $\psi(\mathbf{t}, \mathbf{u}, \mathbf{m}(\mathbf{t}))$ to denote the formula below ($b, b' \in B$): $\psi(\mathbf{t}, \mathbf{u}, (-\infty, b]) \stackrel{\text{def}}{=} \sum_j a_j (x_j + \mathbf{u}(j)) \leq b$; $\psi(\mathbf{t}, \mathbf{u}, [b, +\infty)) \stackrel{\text{def}}{=} \sum_j a_j (x_j + \mathbf{u}(j)) \geq b$ and $\psi(\mathbf{t}, \mathbf{u}, [b, b']) = ((\sum_j a_j (x_j + \mathbf{u}(j)) \leq b') \wedge ((\sum_j a_j (x_j + \mathbf{u}(j)) \geq b))$. We write $\mathbf{G}^*(T, B, U)$ to denote the set of guards of the form $\psi(\mathbf{t}, \mathbf{u}, \mathbf{m}(\mathbf{t}))$ where $\mathbf{t} \in T$, U is the finite set of updates from P and $\mathbf{m} : T \rightarrow I$. Each guard in $\mathbf{G}^*(T, B, U)$ is of linear size in the size of P . We denote $\tilde{\Delta}$ the set of transitions $Q' \times \mathbf{G}^*(T, B, U) \times U \times Q'$. Note that the transitions in $\tilde{\Delta}$ do not contain guards with disjunctions and $\tilde{\Delta}$ is finite. We also define a function proj which associates to $w \in \tilde{\Delta}^\omega$ the ω -sequence $\text{proj}(w) : \mathbb{N} \rightarrow 2^X \times I^T$ such that for all $i \in \mathbb{N}$, if $w(i) = \langle \langle q, \mathbf{m} \rangle, \mathbf{g}, \mathbf{u}, \langle q', \mathbf{m}' \rangle \rangle$ and $\mathbf{l}(q) \cap X = L$ then $\text{proj}(w)(i) \stackrel{\text{def}}{=} \langle L, \mathbf{m} \rangle$.

We show that it is possible to build a finite set Y_P of path schemas over $\tilde{\Delta}$ such that if $P' = p'_1(l'_1)^+ p'_2(l'_2)^+ \dots p'_{k'}(l'_{k'})^\omega$ is a path schema in Y_P and ρ is a run $\langle \langle q_0, \mathbf{m}_0 \rangle, \mathbf{v}_0 \rangle \rightarrow \langle \langle q_1, \mathbf{m}_1 \rangle, \mathbf{v}_1 \rangle \rightarrow \langle \langle q_2, \mathbf{m}_2 \rangle, \mathbf{v}_2 \rangle \dots$ respecting P' we have that $\text{proj}(\text{lab}(\rho)) = \text{ft}(\rho)$. This point will be useful for Algorithm 1. The following theorem lists the main properties of the set Y_P .

Theorem 14. *Given a flat counter system S , a minimal path schema P , a resource $R = \langle X, T, B \rangle$ coherent with P and a configuration $\langle q_0, \mathbf{v}_0 \rangle$, there is a finite set of path schemas Y_P over $\tilde{\Delta}$ satisfying (1)–(6) below.*

1. No path schema in Y_P contains guards with disjunctions in it.
2. There exists a polynomial $q^*(\cdot)$ such that for every $P' \in Y_P$, $\text{len}(P') \leq q^*(\text{len}(P) + \text{card}(T) + \text{card}(B))$.
3. Checking whether a path schema P' over $\tilde{\Delta}$ belongs to Y_P can be done in polynomial time in $\text{size}(P) + \text{card}(T) + \text{card}(B)$.
4. For every run ρ respecting P and starting at $\langle q_0, \mathbf{v}_0 \rangle$, we can find a run ρ' respecting some $P' \in Y_P$ such that $\rho \models \phi$ iff $\rho' \models \phi$ for every ϕ built over R .
5. For every run ρ' respecting some $P' \in Y_P$ with initial values \mathbf{v}_0 , we can find a run ρ respecting P such that $\rho \models \phi$ iff $\rho' \models \phi$ for every ϕ built over R .
6. For every ultimately periodic word $w \cdot u^\omega \in \mathcal{L}(P')$, for every ϕ built over R checking whether $\text{proj}(w \cdot u^\omega), 0 \models_{\text{symbol}} \phi$ can be done in polynomial time in the size of $w \cdot u$ and in the size of ϕ .

5.3 Main Algorithm

In Algorithm 1 below, a polynomial $p^*(\cdot)$ is used. We can define polynomial $p^*(\cdot)$ using the small solutions for constraint systems [2], see details in [6]. Note that

\mathbf{y}' is a refinement of \mathbf{y} (for all i , we have $\mathbf{y}'[i] \approx_{2td(\phi)+5} \mathbf{y}[i]$) in which counter values are taken into account.

Algorithm 1 The main algorithm in NP with inputs $S, c_0 = \langle q, \mathbf{v}_0 \rangle, \phi$

- 1: guess a minimal path schema P of S
 - 2: build a resource $R = \langle X, T, B \rangle$ coherent with P and ϕ
 - 3: guess a valid schema $P' = p_1 l_1^+ p_2 l_2^+ \dots p_k l_k^\omega$ such that $\text{len}(P') \leq q^*(\text{len}(P) + \text{card}(T) + \text{card}(B))$
 - 4: guess $\mathbf{y} \in [1, 2td(\phi) + 5]^{k-1}$; guess $\mathbf{y}' \in [1, 2^{p^*(\text{size}(S)+\text{size}(c_0)+\text{size}(\phi))}]^{k-1}$
 - 5: check that P' belongs to Y_P
 - 6: check that $\text{proj}(p_1 l_1^{\mathbf{y}[1]} p_2 l_2^{\mathbf{y}[2]} \dots l_{k-1}^{\mathbf{y}[k-1]} p_k l_k^\omega), 0 \models_{\text{symb}} \phi$
 - 7: build \mathcal{E} over y_1, \dots, y_{k-1} for P' with initial values \mathbf{v}_0 (obtained from Lemma 12)
 - 8: **for** $i = 1 \rightarrow k - 1$ **do**
 - 9: **if** $\mathbf{y}[i] = 2td(\phi) + 5$ **then** $\psi_i \leftarrow$ “ $y_i \geq 2td(\phi) + 5$ ” **else** $\psi_i \leftarrow$ “ $y_i = \mathbf{y}[i]$ ”
 - 10: **end for**
 - 11: check that $\mathbf{y}' \models \mathcal{E} \wedge \psi_1 \wedge \dots \wedge \psi_{k-1}$
-

Algorithm 1 starts by guessing a path schema P (line 1) and an unfolded path schema $P' = p_1 l_1^+ p_2 l_2^+ \dots p_k l_k^\omega$ (line 3) and check whether P' belongs to Y_P (line 5). It remains to check whether there is a run ρ respecting P' such that $\rho \models \phi$. Suppose there is such a run ρ ; let \mathbf{y} be the unique tuple in $[1, 2td(\phi) + 5]^{k-1}$ such that $\mathbf{y} \approx_{2td(\phi)+5} \text{iter}_{P'}(\rho)$. By Proposition 6, we have $\text{proj}(p_1 l_1^{\mathbf{y}[1]} p_2 l_2^{\mathbf{y}[2]} \dots l_{k-1}^{\mathbf{y}[k-1]} p_k l_k^\omega), 0 \models_{\text{symb}} \phi$. Since the set of tuples of the form $\text{iter}_{P'}(\rho)$ is characterized by a system of equations, by the existence of small solutions from [2], we can assume that $\text{iter}_{P'}(\rho)$ contains only small values. Hence line 4 guesses \mathbf{y} and \mathbf{y}' (corresponding to $\text{iter}_{P'}(\rho)$ with small values). Line 6 precisely checks $\text{proj}(p_1 l_1^{\mathbf{y}[1]} p_2 l_2^{\mathbf{y}[2]} \dots l_{k-1}^{\mathbf{y}[k-1]} p_k l_k^\omega), 0 \models_{\text{symb}} \phi$ whereas line 11 checks whether \mathbf{y}' encodes a run respecting P' with $\mathbf{y}' \approx_{2td(\phi)+5} \mathbf{y}$.

Lemma 15. *Algorithm 1 runs in nondeterministic polynomial time.*

It remains to check that Algorithm 1 is correct.

Lemma 16. *$S, c_0 \models \phi$ iff Algorithm 1 on inputs S, c_0, ϕ has an accepting run.*

In the proof of Lemma 16, we take advantage of all our preliminary results.

Proof. By way of example, we show that if Algorithm 1 on inputs $S, c_0 = \langle q_0, \mathbf{v}_0 \rangle, \phi$ has an accepting computation, then $S, c_0 \models \phi$. This means that there are $P, P', \mathbf{y}, \mathbf{y}'$ that satisfy all the checks. Let $w = p_1 l_1^{\mathbf{y}'[1]} \dots p_{k-1} l_{k-1}^{\mathbf{y}'[k-1]} p_k l_k^\omega$ and $\rho = \langle \langle q_0, \mathbf{m}_0 \rangle, \mathbf{v}_0 \rangle \langle \langle q_1, \mathbf{m}_1 \rangle, \mathbf{x}_1 \rangle \langle \langle q_2, \mathbf{m}_2 \rangle, \mathbf{x}_2 \rangle \dots \in (Q' \times \mathbb{Z}^n)^\omega$ be defined as follows: for every $i \geq 0$, $q_i \stackrel{\text{def}}{=} \pi_1(\text{source}(w(i)))$, and for every $i \geq 1$, we have $\mathbf{x}_i \stackrel{\text{def}}{=} \mathbf{x}_{i-1} + \text{update}(w(i))$. By Lemma 12, since $\mathbf{y}' \models \mathcal{E} \wedge \psi_1 \wedge \dots \wedge \psi_{k-1}$, ρ is a run respecting P' starting at the configuration $\langle \langle q_0, \mathbf{m}_0 \rangle, \mathbf{v}_0 \rangle$. Since $\mathbf{y}' \models \psi_1 \wedge \dots \wedge \psi_{k-1}$ and $\mathbf{y} \models \psi_1 \wedge \dots \wedge \psi_{k-1}$, by Proposition 6, $(\boxtimes) \text{proj}(p_1 l_1^{\mathbf{y}[1]} p_2 l_2^{\mathbf{y}[2]} \dots l_{k-1}^{\mathbf{y}[k-1]} p_k l_k^\omega), 0 \models_{\text{symb}} \phi$, iff $(\boxtimes \boxtimes) \text{proj}(p_1 l_1^{\mathbf{y}'[1]} p_2 l_2^{\mathbf{y}'[2]} \dots l_{k-1}^{\mathbf{y}'[k-1]} p_k l_k^\omega), 0 \models_{\text{symb}} \phi$. Algorithm 1 guarantees that $\text{proj}(p_1 l_1^{\mathbf{y}[1]} p_2 l_2^{\mathbf{y}[2]} \dots l_{k-1}^{\mathbf{y}[k-1]} p_k l_k^\omega), 0 \models_{\text{symb}} \phi$, whence we have $(\boxtimes \boxtimes)$. Since $\text{proj}(p_1 l_1^{\mathbf{y}'[1]} p_2 l_2^{\mathbf{y}'[2]} \dots l_{k-1}^{\mathbf{y}'[k-1]} p_k l_k^\omega) = \text{ft}(\rho)$, by Lemma 13, we deduce that $\rho, 0 \models$

ϕ . By Theorem 14(5), there is an infinite run ρ' , starting at the configuration $\langle q_0, \mathbf{v}_0 \rangle$ and respecting P , such that $\rho', 0 \models \phi$.

For the other direction, see [6]. □

As a corollary, we can state the main result of the paper.

Theorem 17. $\text{MC}(\text{PLTL}[\mathbf{c}], \mathcal{CFS})$ is NP-complete.

6 Conclusion

Classes of Systems	PLTL[\emptyset]	PLTL[\mathbf{c}]	Reachability
\mathcal{KPS}	NP-complete See [16] for \mathbf{X} and \mathbf{U}	—	P _{TIME}
\mathcal{CPS}	NP-complete	NP-complete (Theo. 17)	NP-complete
$\mathcal{KPS}(n)$	P _{TIME} (Theo. 9)	—	P _{TIME}
$\mathcal{CPS}(n), n > 1$??	NP-complete (Lem. 11)	??
$\mathcal{CPS}(1)$	P _{TIME}	P _{TIME}	P _{TIME}
\mathcal{KFS}	NP-complete See [16] for \mathbf{X} and \mathbf{U}	—	P _{TIME}
\mathcal{CFS}	NP-complete	NP-complete (Theo. 17)	NP-complete

We have investigated the computational complexity of the model-checking problem for flat counter systems with formulae from an enriched version of LTL. Our main result is the NP-completeness of $\text{MC}(\text{PLTL}[\mathbf{c}], \mathcal{CFS})$, significantly improving the complexity upper bound from [7]. This also improves the results about the effective semilinearity of the reachability relations for such flat counter systems from [5,11] and it extends the recent result on the NP-completeness of model-checking flat Kripke structures with LTL from [16] by adding counters and past-time operators. Our main results are presented above and compared to the reachability problem.

As far as the proof technique is concerned, the NP upper bound is obtained as a combination of a general stuttering property for LTL with past-time operators (a result extending what is done in [17] with past-time operators) and the use of small integer solutions for quantifier-free Presburger formulae [2]. There are several related problems which are not addressed in the paper. For instance, the extension of the model-checking problem to full CTL* is known to be decidable [7] but the characterization of its exact complexity is open.

References

1. B. Boigelot. *Symbolic methods for exploring infinite state spaces*. PhD thesis, Université de Liège, 1998.
2. I. Borosh and L. Treybig. Bounds on positive integral solutions of linear Diophantine equations. *American Mathematical Society*, 55:299–304, 1976.
3. M. Bozga, R. Iosif, and F. Konečný. Fast acceleration of ultimately periodic relations. In *CAV'10*, volume 6174 of *LNCS*, pages 227–242. Springer, 2009.

4. H. Comon and V. Cortier. Flatness is not a weakness. In *CSL'00*, volume 1862 of *LNCS*, pages 262–276. Springer, 2000.
5. H. Comon and Y. Jurski. Multiple counter automata, safety analysis and PA. In *CAV'98*, volume 1427 of *LNCS*, pages 268–279. Springer, 1998.
6. S. Demri, A. Dhar, and A. Sangnier. Taming past LTL and flat counter systems. Technical report, 2012. ArXiv.
7. S. Demri, A. Finkel, V. Goranko, and G. van Drimmelen. Model-checking CTL* over flat Presburger counter systems. *JANCL*, 20(4):313–344, 2010.
8. C. Dixon, M. Fisher, and B. Konev. Temporal logic with capacity constraints. In *FRODOS'07*, volume 4720 of *LNCS*, pages 163–177. Springer, 2007.
9. J. Esparza, A. Finkel, and R. Mayr. On the verification of broadcast protocols. In *LICS'99*, pages 352–359, 1999.
10. K. Etessami and T. Wilke. An until hierarchy and other applications of an Ehrenfeucht-Fraïssé game for temporal logic. *I&C*, 160(1–2):88–108, 2000.
11. A. Finkel and J. Leroux. How to compose Presburger accelerations: Applications to broadcast protocols. In *FST&TCS'02*, volume 2256 of *LNCS*, pages 145–156. Springer, 2002.
12. A. Finkel, E. Lozes, and A. Sangnier. Towards model-checking programs with lists. In *Infinity in Logic & Computation*, volume 5489 of *LNAI*, pages 56–82. Springer, 2009.
13. D. Gabbay. The declarative past and imperative future. In *Temporal Logic in Specification, Altrincham, UK*, volume 398 of *LNCS*, pages 409–448. Springer, 1987.
14. C. Haase, S. Kreutzer, J. Ouaknine, and J. Worrell. Reachability in succinct and parametric one-counter automata. In *CONCUR'09*, volume 5710 of *LNCS*, pages 369–383. Springer, 2009.
15. L. Kuhtz. *Model Checking Finite Paths and Trees*. PhD thesis, Universität des Saarlandes, 2010.
16. L. Kuhtz and B. Finkbeiner. Weak Kripke structures and LTL. In *CONCUR'11*, volume 6901 of *LNCS*, pages 419–433. Springer, 2011.
17. A. Kučera and J. Strejček. The stuttering principle revisited. *Acta Informatica*, 41(7–8):415–434, 2005.
18. F. Laroussinie, N. Markey, and P. Schnoebelen. Temporal logic with forgettable past. In *LICS'02*, pages 383–392. IEEE, 2002.
19. F. Laroussinie and P. Schnoebelen. Specification in CTL + past for verification in CTL. *I&C*, 156:236–263, 2000.
20. J. Leroux and G. Sutre. Flat counter systems are everywhere! In *ATVA'05*, volume 3707 of *LNCS*, pages 489–503. Springer, 2005.
21. C. Lutz. NEXPTIME-complete description logics with concrete domains. In *IJ-CAR'01*, volume 2083 of *LNCS*, pages 46–60. Springer, 2001.
22. M. Minsky. *Computation, Finite and Infinite Machines*. Prentice Hall, 1967.
23. D. Peled and T. Wilke. Stutter-invariant temporal properties are expressible without the next-time operator. *IPL*, 63:243–246, 1997.
24. C. Rackoff. The covering and boundedness problems for vector addition systems. *TCS*, 6(2):223–231, 1978.